3

Socially Shared Cognition at Work

Transactive Memory and Group Performance

RICHARD L. MORELAND
LINDA ARGOTE
RANJANI KRISHNAN

A recent review by Moreland, Hogg, and Hains (1994) of social psychological research revealed a resurgence of interest in studying small groups. Much of that resurgence was due to new theoretical approaches for analyzing group behavior. One approach that was especially influential involved analyses incorporating social cognition. Fiske and Goodwin (Introduction, this volume) have argued that valuable insights into many group phenomena can occur when the cognitive processes associated with group membership are studied. Similar claims have been made by other observers as well (e.g., Goethals, 1987; Messick & Mackie, 1989).

There are two general ways in which such cognitive processes could be studied (Ickes & Gonzalez, Chapter 12, this volume). Most researchers focus on cognition about groups, studying such phenomena as self-categorization (Brewer, 1991), outgroup homogeneity effects (Judd, Ryan, & Park, 1991), the group attribution error (Allison &

Messick, 1985), and so on. Their work explores how *individuals* process information about groups and their members (including themselves). A few researchers, however, focus on more complex phenomena involving cognition by groups. Their work explores how *groups* process information for their members. Information processing by groups requires *socially shared cognition,* that is, collaboration among members who seek to encode, interpret, and recall information together rather than apart. Although little is yet known about socially shared cognition, it has captured considerable attention (see Klimoski & Mohammed, 1994; Larson & Christensen, 1993; Resnick, Levine, & Teasley, 1991) and seems likely to generate much new research.

WORK GROUP CULTURES

One intriguing phenomenon that arises from socially shared cognition is culture. Psychologists are accustomed to analyzing culture at the societal level (Markus & Kitayama, 1991; Triandis, 1989) or within large organizations (Schein, 1990; Trice & Beyer, 1984), but small groups can develop cultures as well. Evidence for such cultures can be found in research on families (Wolin & Bennett, 1984), sports teams (Fine, 1979), and work groups (Rentsch, 1990). Cultural analyses of work groups have been especially popular. In a recent review of this research, Levine and Moreland (1991) suggested that every group of workers develops its own unique culture. Of course, the culture of a particular group may be weak, depending on such factors as the age of the group, the homogeneity and stability of its composition, the levels of process and outcome interdependence among members, and the group's cohesion. And even if a work group's culture is strong, it may be tacit, sensitive, or dynamic and thus difficult to study. Nevertheless, cultural analyses of work groups are possible and can reveal aspects of group dynamics that might otherwise be overlooked.

What, exactly, is the "culture" of a work group? According to Levine and Moreland (1991), it consists of two related elements, namely socially shared knowledge and a set of customs. Cultural knowledge can focus on the group, its members, or the work that

Knowledge About the Group

- What are our norms?
- How successful are we?
- Who are our allies/enemies?

Knowledge About Group Members

- How do I fit in?
- What are the cliques?
- Who is good at what?

Knowledge About Work

- What are our working conditions?
- Why do we work?
- How is performance evaluated?

Figure 3.1. Cultural Knowledge in Work Groups: Examples of Questions for Which Workers Seek Common Answers

they perform (cf. Cannon-Bowers, 1993). Some examples of knowledge related to each of these foci are provided in Figure 3.1. The examples are phrased as questions that often concern a work group's members. Their concern is what drives the development of culture, which provides group members with the answers to such questions. Customs, which include routines (Gersick & Hackman, 1990), accounts (Orr, 1990), jargon (Truzzi & Easto, 1972), rituals (Vaught & Smith, 1980), and symbols (Riemer, 1977), embody this cultural knowledge and thus serve (often in subtle ways) to communicate and validate it.

Although the questions shown in Figure 3.1 may seem simple, answering them often involves complex cognitive/social processes. Consider, for example, how workers evaluate their group's success. Success implies some standard, which may be objective (e.g., meeting production goals, avoiding accidents) or subjective. Subjective standards include temporal and social comparisons (Levine & Moreland, 1987); the latter comparisons may focus on various outgroups. To evaluate their group's success, members must thus agree about which

standards to use and how they should be applied. It may also be necessary to reconcile their own evaluation of the group with evaluations by outsiders. Even after a final evaluation is made, other questions may soon arise, such as why the group is succeeding or failing (Leary & Forsyth, 1987) and how it might be improved. Discovering the cliques within a group can also be difficult for workers (Krackhardt, 1990), who may be reluctant to provide one another with direct information about who they like or dislike within the group. Indirect information about cliques can be gathered through observations of social behavior, but members must agree about which behaviors are important and how they should be interpreted. And once again, further questions are likely to arise, such as how certain people became friends or enemies and whether their feelings toward one another are likely to change. Finally, workers may also struggle to explain why they work. Intrinsic and extrinsic motivators must be identified, assessed, weighted, and combined. These activities often require considerable discussion among group members (Salancik & Pfeffer, 1978).

Our own research has focused on a question asked by many members of work groups: Who is good at what? Workers are not clones—every member offers the group a unique set of abilities, skills, and knowledge that may be relevant to a variety of tasks. Although self-deception can occur, many workers seem to have an accurate sense of their own competence (Shrauger & Osberg, 1981). Evaluating the competence of other group members, however, can be difficult. Claims of personal competence by coworkers cannot always be trusted because they may reflect efforts at impression management (Gardner, 1992). Accepting coworkers' evaluations of one another's competence can be risky as well because these "secondhand" evaluations are often based on limited information (Gilovich, 1987) and may reflect more subtle impression management efforts by the people who provide them (Cialdini, 1989). The best approach for evaluating competence may be to observe how well each group member actually performs various tasks. But even this method has its drawbacks. For example, opportunities for displaying some kinds of abilities, skills, or knowledge may be rare, and even when they do occur, such opportunities may not be distributed equally among group members. Various biases, such as racism or

sexism, may also distort evaluations of competence, causing some coworkers to seem better or worse than they really are (Ridgeway, 1991).

Given these and other problems regarding evaluations of competence, why should the members of a work group worry about who is good at what? We believe that such knowledge is valuable because it can improve the group's performance in several ways. First, when group members know who is good at what, they can plan their work more sensibly, assigning tasks to the people who will perform them best. As a result, the group can make optimal use of its human resources. Second, coordination among workers is likely to improve when they know who is good at what. Familiarity with one another's strengths and weaknesses helps workers to *anticipate* behavior rather than merely *react* to it. Even when tasks have not been assigned to particular people, workers can thus interact more smoothly and efficiently. Finally, unexpected problems can be solved more quickly and easily when workers know who is good at what (Moreland & Levine, 1992a). Such knowledge allows workers to match problems with the people most likely to solve them. Those people can then be asked for help, or the problems can simply be turned over to them.

If these claims seem unconvincing, then just imagine a work group whose members knew nothing about who was good at what. In such a group, sensible planning would be impossible. Tasks might be assigned to workers at random or, worse yet, on the basis of such irrelevant cues as appearance or demeanor. Or the group might fail to plan its work at all, simply allowing people to work on whatever tasks they liked best (cf. Hackman & Morris, 1975; Weingart, 1992). It would be lucky indeed if any of these options led to optimal use of the group's resources. Coordination would probably suffer as well, as confused workers struggle to make sense out of one another's behavior, and problems might become more troublesome, as unqualified workers try (or are asked) to solve them.

Do work groups really perform better when their members know who is good at what? The available evidence (e.g., Bottger, 1984; Libby, Trotman, & Zimmer, 1987; Littlepage & Silbiger, 1992; Stasser, Stewart, & Wittenbaum, 1995; Yetton & Bottger, 1982), drawn mostly from laboratory research on the recognition of expertise

in decision-making groups, suggests that they do. Libby and his colleagues (Libby, Trotman, & Zimmer, 1987), for example, asked groups of bank loan officers to review the financial profiles of various companies and predict whether each company would go bankrupt. The subjects performed this task twice, first individually and then in groups. The companies were real (but disguised), so their actual outcomes were known, allowing the accuracy of all these predictions to be assessed. The ability of groups to recognize expertise was measured in two ways. First, each group was asked, after all the predictions were made, to identify its best member. The relative accuracy of that person's individual predictions, versus those of the other group members, was then evaluated to see how much better than his colleagues he really was. Second, the relationship between each person's relative accuracy and his influence on the group (how often he and the group made the same predictions) was evaluated to see how much better than his colleagues the most influential person was. Analyses of group performance showed that both diversity in group members' expertise and the ability of groups to recognize expertise were associated with more accurate group predictions.

Littlepage and Silbiger (1992) asked groups of college students to take trivia tests on such topics as geography, sports, and entertainment. After answering each test question, groups were asked to assign it a point value, indicating their confidence in the answer. It was thus possible to measure not only the actual expertise of a group (the number of questions answered correctly) but also its ability to recognize expertise. That ability was measured by first calculating the total number of points a group earned and then dividing that sum by the total number of points that could have been earned if the group had assigned the largest possible values to the questions that it answered correctly. Analyses of these data showed that the ability to recognize expertise was associated with better group performance, even when actual expertise was taken into account.

Finally, Stasser et al. (1995) provided groups of college students with clues to a hypothetical murder and then asked them to discover who (among several suspects) committed the crime. Within each group, some clues were distributed to everyone, whereas other clues were distributed in ways that provided every member with expertise on a different suspect. The latter clues were critical for solving the crime. Every group was informed that its members might receive different sets of clues. Information about exactly who knew what, however, varied from one group to another. In some groups, but not others, each member was privately informed about his or her own area of expertise. And in some groups, but not others, everyone's area of expertise was publicly announced. Simply informing group members about their own expertise had little impact on group performance, but announcing everyone's expertise helped group members to discuss more of the critical clues and thereby solve the crime.

TRANSACTIVE MEMORY SYSTEMS

A key factor in the performance of many work groups is the ability to remember information relevant to their tasks. Once again, knowing who is good at what, or more narrowly, who *knows* what, could be helpful. Although memory is usually viewed as a cognitive process that occurs within individual minds, some psychologists (e.g., Clark & Stephenson, 1989; Walsh & Ungson, 1991; Wegner, 1987) have become interested in the cognitive/social process of collective memory. One of the most intriguing analyses of such memory was offered by Wegner (1987). Wegner noted that few people rely exclusively on their own memories, which are limited and often faulty. Instead, most people supplement their own memories with external aids. These memory aids include both objects (e.g., address and appointment books, lists of uncompleted projects) and other people (e.g., friends, spouses, coworkers). Wegner was especially interested in the use of people as memory aids. He argued that in many groups, a *transactive memory system* develops for the purpose of ensuring that important information is not forgotten. Such a system combines the knowledge possessed by particular group members with a shared awareness of who knows what. So when group members need information, but either cannot recall it themselves or are uncertain about the accuracy of their own memories, they can turn to one another for

help. A transactive memory system thus provides access to more and better information than any single group member could remember alone. Although research on transactive memory systems is scarce and limited to couples (see Wegner, 1987; Wegner, Erber, & Raymond, 1991; Wegner, Giuliano, & Hertel, 1985), the available evidence suggests that such systems may be both common and useful in groups.

Transactive memory systems should be very helpful for work groups, especially groups that perform complex tasks requiring considerable knowledge. When are such systems likely to develop? If transactive memory systems are viewed as cultural phenomena, then many of the factors associated with stronger group cultures may be important. These factors, which were mentioned earlier, include the age of a work group, the homogeneity and stability of its composition, the levels of process and outcome interdependence among workers, and the group's cohesion. Our research has focused on the role of *shared experience* (which incorporates several of these factors) in both the development and operation of transactive memory systems. As people work together on group tasks, they not only acquire more information about those tasks themselves, but also discover whatever relevant information other group members possess. When these two kinds of information are combined, a transactive memory system becomes available for use by the group.

If transactive memory systems are helpful to work groups, and shared experience is indeed important for the development of those systems, then a group's performance should improve as its members spend more time together working at similar or related tasks. This claim has two general implications for the management of work groups. First, managers should try to minimize changes in a group's composition (due to absenteeism, leaves, transfers, or turnover) because the arrival of newcomers and/or departure of oldtimers from the group can seriously disrupt its transactive memory system. Although no one has studied this phenomenon directly, several studies suggest that groups do learn to make better use of their members' knowledge over time (e.g., Larson, Foster-Fishman, & Keys, 1994; Watson, Kumar, & Michaelsen, 1993; Watson, Michaelson, & Sharp, 1991), especially if their composition remains stable.

Argote (1993), for example, reviewed research on group learning and found some evidence of "learning curves" similar to those displayed by individuals. Over time, a group's task performance improves, but at a decreasing rate (Leavitt, 1951; Perlmutter & De Montmollin, 1952). Argote argued that group learning may reflect learning by individual group members, changes in the group's structure or dynamics that improve productivity, or the adoption of useful methods developed by other groups to perform the same task. This analysis suggests that turnover should disrupt group learning because (a) newcomers generally know less than oldtimers about performing the group's task and (b) newcomers and oldtimers are often unfamiliar with one another's abilities and interests. A recent experiment by Argote, Insko, Yovetich, and Romero (1995), in which laboratory groups spent several consecutive work periods producing origami objects of varying complexity, showed that turnover indeed disrupts group learning and that its impact grows worse as time passes. Moreover, training newcomers in origami skills before they entered those groups did not weaken that impact. This finding suggests that turnover was harmful primarily because newcomers and oldtimers were unfamiliar with one another and thus had difficulty working together effectively. A similar experiment by Devadas and Argote (1995), in which much of the damage caused by turnover was averted by strengthening group structure (assigning specific roles to members and prescribing work procedures), suggests the same conclusion.

Research by Goodman and his colleagues (see Goodman & Shah, 1992) on how "familiarity" affects workers in coal-mining crews also seems relevant to this issue. In coal mines, familiarity might involve equipment, terrain, or personnel. Changes in any of these factors should decrease familiarity and thus be harmful. This is just what Goodman and his colleagues found after analyzing archival data from several mines. In one study (Goodman & Garber, 1988), familiarity was associated with fewer accidents among pairs of crew members who worked very closely together (e.g., roof bolters and bolter helpers). Although familiarity with terrain had somewhat more impact than personnel familiarity, the latter factor was clearly important, especially when groups worked in less familiar terrain. In a later

study (Goodman & Leyden, 1991), familiarity was associated with higher levels of crew productivity, even after other labor, technology, and environment factors were taken into account. Once again, personnel familiarity had less impact than familiarity with terrain—perhaps crews adjust more easily to changes in personnel than to changes in terrain. However, changes in personnel were clearly harmful, and the fact that some of the effects of familiarity on productivity were mediated by labor (but not technology or environment) factors provided further evidence for the importance of such changes.

Our claim about the benefits of shared experience has another general implication for the management of work groups, namely that managers should train the members of such groups together rather than apart. Most organizations emphasize individual learning in their training programs. Participants in those programs work on their own, under the guidance of instructors, to learn how various tasks should be performed. The fact that they may later perform those tasks in groups is largely ignored. Some training programs do include group activities, of course, but these often focus on general topics such as leadership or cooperation (Oddou, 1987; Silberman, 1990; Tetrault, Schriesheim, & Neider, 1988) that are relevant to any group. Specific knowledge about the actual groups that workers will later form or join is seldom provided. In particular, few workers learn much about who knows what within those groups. As a result, transactive memory systems develop more slowly and operate less efficiently than they could if group members were trained together. Training the members of a work group together rather than apart would not only provide each person with the information needed to perform tasks well, but also help him or her to learn what everyone else in the group knows about those tasks.

Some organizations have recently begun to train work group members together rather than apart. General Motors, at its Saturn automobile manufacturing plant, has received considerable publicity in this regard, but other examples can also be found, such as the U.S. Army's unit personnel replacement system (Griffith, 1989) and efforts by several airlines to provide team training for cockpit crews (Oberle, 1990). Although enthusiastic claims have been made about

these innovative training programs, convincing evidence regarding their benefits is scarce. Few researchers have actually compared the performance of work groups whose members were trained together versus apart (see Dyer, 1985, for a review), and the results of those studies are mixed. Some studies (e.g., George, 1967) have shown that group training is superior to individual training, whereas others (e.g., Briggs & Naylor, 1965) have shown that group training is inferior. Several studies (e.g., Laughlin & Sweeney, 1977) showed no difference at all between the two forms of training. These divergent findings are puzzling, and their interpretation is further clouded by methodological problems and a general lack of theory. Clearly, further research is needed to determine whether group or individual training yields better group performance.

We have embarked on a program of research designed to explore this issue. Our research involves laboratory experiments in which groups of subjects are trained to perform a complex task. Various types of training, focusing on either groups or individuals, are provided and their effects on group performance are compared. Several hypotheses are tested, but our two general predictions are that (a) group performance improves when group members are trained together rather than apart and (b) the benefits of group training are due to the development and operation of transactive memory systems.

OUR FIRST EXPERIMENT

Our first experiment (Liang, Moreland, & Argote, 1995) was performed at Carnegie Mellon University, where 90 students enrolled in undergraduate business courses served as subjects. These subjects were randomly assigned to small work groups, each containing three persons of the same sex. We chose a task that simulated the type of work found in many manufacturing organizations. Every group was asked to assemble the AM portion of an AM/FM radio using materials from kits. Each kit included a circuit board and dozens of mechanical and electronic components. The circuit board contained prepunched holes marked with special symbols indicating where different components

should be placed. To assemble a radio, the subjects had to insert each component into the circuit board at the proper place and then wire all of the components together in the proper manner.

Two types of training were provided for this difficult task. Half of the groups were randomly assigned to a group training condition, whereas the other half were assigned to an individual training condition. The training *format* differed across these conditions. In the group training condition, members of the same group were trained together, whereas in the individual training condition, they were trained apart. The *content* of training, however, did not differ across conditions—the same trainer presented the same information to everyone in similar ways. If group performance on a task like this depends solely on the knowledge of individual group members, then training them together or apart should not matter. But if group performance depends partly on transactive memory, which allows group members to use one another's knowledge as well, then training them together should be advantageous, because it helps a transactive memory system to develop.

The experiment was carried out in two sessions, each lasting for about an hour. The first session focused on training, whereas the second session (which occurred about one week later) focused on testing. In the group training condition, the members of each group participated in both sessions together. But in the individual training condition, each person participated in his or her own training session, and group members did not meet or work together until the testing session.

When subjects arrived for the training session, they were told that our research examined how training affects work group performance. An overview of the experiment was then provided so that everyone knew what to expect. Every subject realized that at the next (testing) session, he or she would be working in a group whose performance would be evaluated. Subjects in the group training condition expected to remain in their current work group, but subjects in the individual training condition did not know who their coworkers might be. To provide an incentive for good performance, cash awards were offered for members of the best work groups.

The actual training began with a demonstration by the experimenter of how the radio's components should be inserted into the circuit board and wired together. Subjects were allowed to ask questions during this demonstration, which lasted for about 15 minutes. They were then given up to 30 minutes to assemble a single radio themselves, again asking questions if necessary. Finally, the experimenter examined the radio carefully, identified any errors that the subjects made, and explained how those errors could be corrected.

When groups arrived for the testing session, they were first given a memory test. Each group had 7 minutes to recall (without access to any materials or instructions) how a radio should be assembled. Group members collaborated at this task and recorded whatever they could remember together on a single sheet of paper. Each group was then given up to 30 minutes to actually assemble a radio, without consulting its recall sheet or receiving any help from the experimenter. The subjects were told to work as quickly as possible, but also to make as few errors as possible. Every group's activities were recorded on videotape. Finally, each subject completed a brief questionnaire that provided background information about that person and measured his or her beliefs about the group and its task.

We analyzed three measures of group performance, namely (a) how well each group recalled the procedure for assembling a radio, (b) how quickly each group assembled its radio, and (c) how many assembly errors that radio contained. Groups in the two training conditions assembled their radios at about the same speed. However, there were significant differences between conditions in both procedural recall and assembly errors. These differences are shown in Figure 3.2. As we expected, groups whose members were trained together rather than apart recalled more about how a radio should be assembled, and produced radios containing fewer assembly errors.

Videotapes of the groups allowed us to explore several factors that could have produced these effects. Two judges, one of whom was blind to the research hypotheses and to each group's condition, were given a list of specific behaviors exemplifying each factor. They were then asked to watch each videotape carefully, keeping those behaviors in mind, and make an overall rating of the group on each factor. Intraclass correlations, computed for each factor, indicated that these ratings were made reliably.

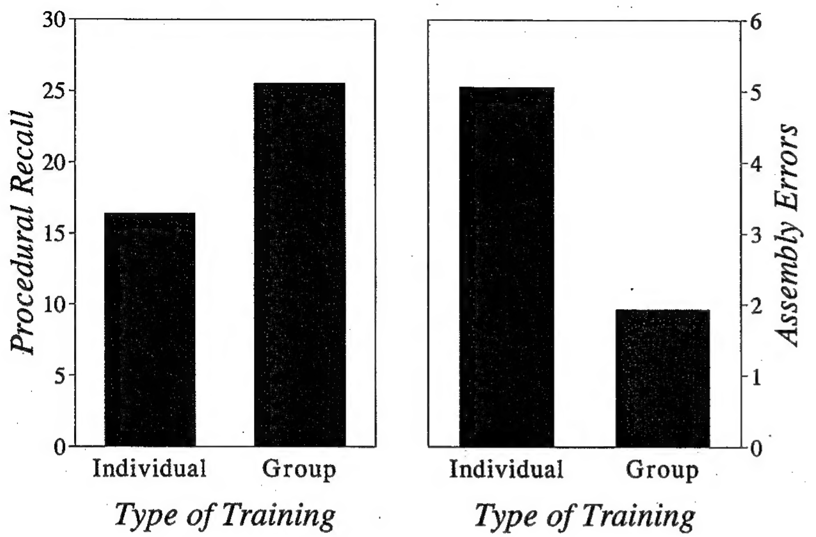

Figure 3.2. Effects of Group and Individual Training on Procedural Recall (Left) and Assembly Errors (Right) in the First Experiment

We were especially interested in three factors that may reflect the operation of transactive memory systems. The first factor was *memory differentiation,* or the tendency for group members to specialize in recalling different aspects of assembling the radio. One person, for example, might remember where certain components should be inserted into the circuit board, whereas another one remembered how those components should be wired together. The second factor was *task coordination,* or the ability of group members to work together smoothly while assembling the radio. In groups with stronger transactive memory systems, there should be less confusion, fewer misunderstandings, and greater cooperation. Finally, the third factor was *task credibility,* or how much group members trusted one another's knowledge about the radio. In groups with stronger transactive memory systems, there should be less need to claim expertise, better acceptance of any procedural suggestions, and less criticism of work by others. Scores on these three factors were highly correlated with one another,

as might be expected if they measured the same underlying phenomenon. We thus created a composite index of transactive memory activity by averaging together each group's scores on these factors.

Three other factors were also coded from the videotapes. These factors seemed relevant to group performance, and might have varied across training conditions, but were not assumed to measure any single underlying phenomenon. The first factor was *task motivation,* or how eager group members were to win the prize by assembling their radio quickly and correctly. Group members whose motivation is stronger should express more enthusiasm for the task, encourage one another more often, and work harder. The second factor was *group cohesion,* or the level of interpersonal attraction among group members. Members of more cohesive groups should sit closer together, speak more warmly to one another, and so on. Finally, the third factor was *social identity,* or the tendency for subjects to think about themselves as group members rather than individuals. This was the only factor not evaluated through ratings. Instead, the judges counted how often personal pronouns (e.g., "I") and collective pronouns (e.g., "We") were used while the members of each group assembled their radio. The ratio of collective pronouns to all pronouns used was then computed and served as a measure of social identity (cf. Cialdini et al., 1976; Veroff, Sutherland, Chadiha, & Ortega, 1993).

We expected stronger transactive memory systems to develop in groups whose members were trained together rather than apart, and that is what we found. Scores on all three transactive memory factors (memory differentiation, task coordination, task credibility) and on the composite index of transactive memory were significantly higher in the group training condition. Only one of the other three factors differed significantly across training conditions. Whether group members were trained together or apart had little impact on task motivation or group cohesion scores, but higher social identity scores were earned by groups whose members were trained together.

Group training clearly produced better group performance and stronger transactive memory systems. Are those findings connected? We believed that the effects of group training on group performance were *mediated* by transactive memory systems. To explore that issue,

a series of multiple regression analyses was performed (Baron & Kenny, 1986). The goal of those analyses, as shown in Figure 3.3, was to separate the direct and indirect (mediated) effects of training methods on group performance. If the direct effects (Path C) are weak while the indirect effects (Paths A and B) are strong, then that would provide evidence for mediation. Our measure of group performance for these analyses was assembly errors—training methods did not affect how quickly the groups assembled their radios, and the effects of training methods on procedural recall occurred before transactive memory was measured. We began by regressing assembly errors on training methods (coded as a dummy variable). The overall effects of training methods were significant and accounted for 42% of the variance. We then regressed transactive memory (using scores from the composite index) on training methods. The effects of training methods on transactive memory (Path A) were also significant and accounted for 70% of the variance. Finally, we regressed assembly errors on training methods and transactive memory simultaneously. The effects of transactive memory on assembly errors (Path B) were significant, but the effects of training methods (Path C) were not. Together, training methods and transactive memory accounted for 57% of the variance in assembly errors.

These results suggest that the impact of group training on group performance was indeed mediated by transactive memory systems, because when differences among groups in the strength of those systems were taken into account, training methods no longer mattered. Although we did not expect social identity to play a similar role, analogous regression analyses involving that factor were also performed. The results provided no evidence of mediation—when differences among groups in social identity were taken into account, the effects of training methods on group performance remained significant.

Our first experiment provided clear evidence that a work group's performance can be improved by training its members together rather than apart. As we expected, groups whose members were trained together recalled more about how to assemble their radios and made fewer errors while assembling those radios. We also expected and found stronger transactive memory systems in groups whose members

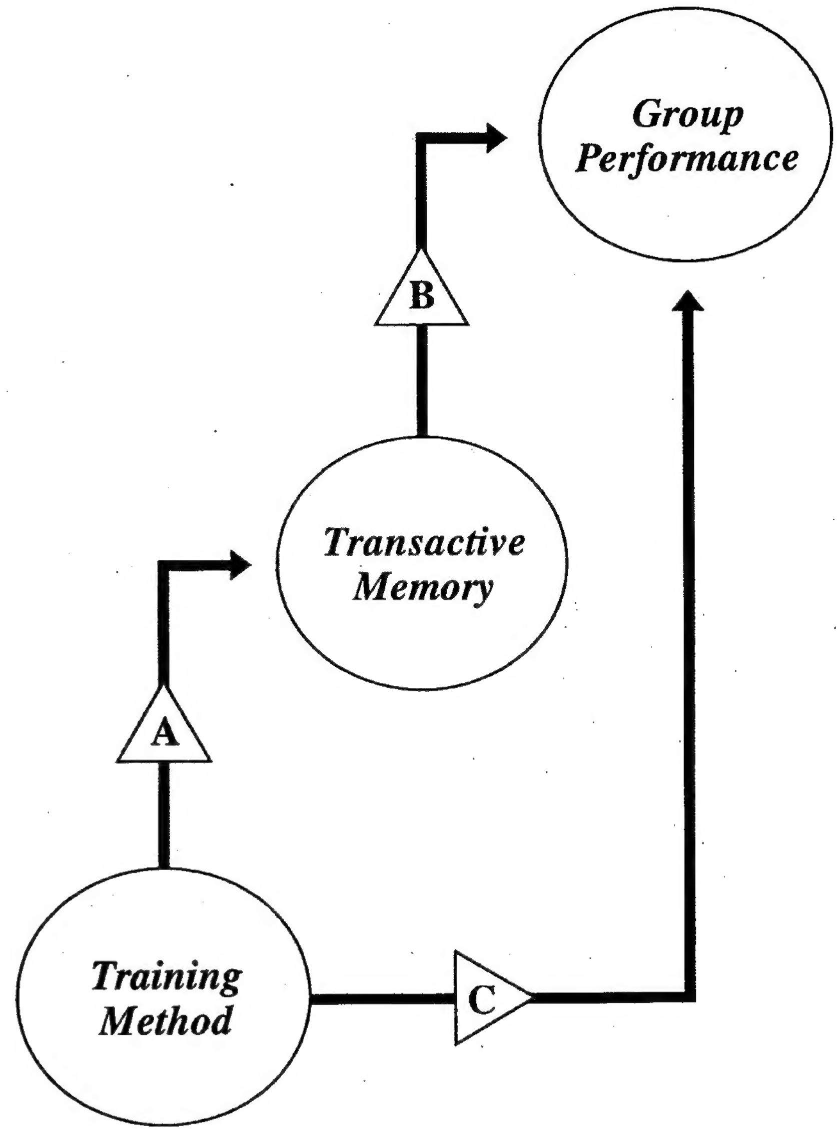

Figure 3.3. Testing Whether Transactive Memory Mediates the Effects of Training Methods on Group Performance

were trained together. The members of those groups, while working on their radios, were more likely to (a) recall different aspects of the procedure, (b) coordinate their activities, and (c) trust one another's expertise. These findings represent the first evidence that transactive memory systems can develop in work groups, as Wegner (1987) claimed. Finally, we found that transactive memory mediates the effects of group training on group performance. When differences in the strength of their transactive memory systems are taken into account, groups whose members were trained together perform no better than groups whose members were trained apart.

OUR SECOND EXPERIMENT

Although these initial results were encouraging, several important issues remained unresolved. A second experiment (Moreland & Wingert, 1995) was thus performed. The subjects for this experiment were 186 students enrolled in undergraduate psychology courses at the University of Pittsburgh. Many of the procedures for this experiment were identical to those described earlier. Once again, subjects were randomly assigned to small work groups, each containing three persons of the same sex. These groups were then randomly assigned to group or individual training conditions. The task and materials remained the same, and each group again participated in training and testing sessions, which lasted for about an hour each and were held one week apart. The training sessions for this experiment were changed slightly for groups in the individual training conditions. Rather than participating in separate training sessions, the members of these groups were trained in the same room at the same time but were not allowed to talk with one another or to observe one another while assembling the radios. This simplified the experiment and made the experiences of subjects in the group and individual training conditions more comparable. Another minor procedural change, which affected groups in every training condition, involved the testing sessions. At the beginning of those sessions, all of the subjects completed brief questionnaires (to be described later in more detail)

that measured feelings and thoughts about their groups. The general format (including time limits) for both the training and testing sessions was otherwise unchanged. During the training sessions, group members again watched a demonstration of how to assemble a radio and then practiced assembling one themselves with guidance and feedback from the experimenter. During the testing sessions, groups again recalled how to assemble a radio and then assembled one entirely on their own, working as quickly and accurately as possible. All of the groups were again videotaped while they worked.

One simple, but important, goal for this experiment was to replicate our initial results. The same individual and group training conditions were thus re-created so that their effects on group performance could be reexamined. Our first experiment showed that group training led to better performance and that this advantage was due to the development of transactive memory systems. Similar results were expected from our second experiment.

A second goal for this experiment was to evaluate alternative explanations for the improvement in group performance associated with group training. For example, group training allows group development to occur. Newly formed groups suffer from many problems (e.g., anxiety about acceptance, interpersonal conflicts, uncertainty about norms) that can impair their performance (LaCoursiere, 1980; Tuckman, 1965). Training the members of a work group together provides them with extra time and opportunities to solve such problems. Perhaps the performance advantages of group training are due more to group development than to transactive memory. Several results from the first experiment challenged this conclusion—neither group cohesion nor social identity (both potential indexes of group development) mediated the effects of group training on group performance, and although questionnaire responses revealed that the members of a few groups were acquainted before the experiment began, those groups (which were probably more developed) did not perform especially well. Nevertheless, the role of group development in group training seemed to deserve a closer examination.

Strategic learning is another alternative explanation for improved performance by work groups whose members were trained

together rather than apart. Working together creates coordination problems, some of which can be solved through simple, generic strategies that do not require much information about who knows what. These strategies include building commitment to the group and its task, organizing task activities, and maintaining harmony among group members. Perhaps the performance advantages of group training are due more to the acquisition of these strategies than to transactive memory. What people may learn during such training is how to perform the task well in any group, rather than in a specific group. None of the results from the first experiment really challenged this conclusion, so the role of strategic learning in group training clearly deserved a closer examination as well.

To evaluate the contributions of group development and strategic learning to the performance advantages of group training, we added two new conditions to our second experiment. One new condition was identical to the individual training condition, except for the addition of a brief team-building exercise that was designed to encourage group development. This exercise, adapted from McGrath (1993), occurred at the end of the training session after subjects finished working on their radios. Group members were seated at a round table and given 10 minutes to develop a quiz that could be used by the university to evaluate juniors and seniors who wanted to become "mentors" for freshmen during fall orientation sessions. Six multiple-choice questions were required, two on each of three topic areas (history and traditions, rules and regulations, locations of buildings and services). Three alternative answers were required for each question, with the correct answer marked on the quiz. The other new condition was identical to the group training condition, except that all of the groups were scrambled between their training and testing sessions. That is, the subjects were reassigned to new groups in ways that separated people who were trained together. When each of these groups began its testing session, its members were thus strangers to one another. The subjects were not told that this scrambling would occur until the end of their training sessions.

The team-building condition was meant to encourage group development, without providing subjects the information (who knows

what?) needed to produce a transactive memory system. The reassignment condition was meant to disable existing transactive memory systems by making them irrelevant, leaving strategic learning as the primary benefit of group training. Insofar as group development and strategic learning each contribute to the performance advantages of group training, performance by groups in these two new conditions should be good. But what if transactive memory systems mediate the effects of group training on group performance, as the results from our first experiment indicated? Groups in the team-building condition lack transactive memory systems, whereas groups in the reassignment condition have transactive memory systems that are no longer useful. As a result, the performance of the groups in these two new conditions should be poor.

A final goal for this experiment was to explore the impact of turnover on transactive memory systems. One benefit of such systems is that no one needs to know everything—each person can rely on the knowledge of others to some extent, so long as everyone agrees about who knows what. But what if someone leaves the group, taking away some valuable knowledge that no one else possesses? Other members of the group may feel regret, because they could have acquired that knowledge themselves, but chose not to do so, and frustration, because they know where that knowledge resides, but no longer have access to it. And, of course, the group's performance may suffer on any tasks requiring that knowledge. In our experiment, these problems could arise among groups in the reassignment condition. Those groups experienced sudden and dramatic turnover after their training sessions. When the members of those groups later began their testing sessions, working with strangers whose knowledge about assembling radios was largely unknown, some of them probably wished they had learned more about the task, wanted to recontact old coworkers and ask task-related questions, or wondered whether such questions could be answered by their new coworkers. How harmful is all this for group performance? We hoped that groups in the reassignment condition might help to answer that question.

Our first data analyses focused on the brief questionnaire that subjects completed at the beginning of their testing sessions. The

purpose of that questionnaire was to evaluate how the various training conditions affected subjects' feelings and thoughts about their groups. The questionnaire contained 10 questions that subjects answered by making ratings on 7-point scales. Some of the questions (e.g., "Does this work group seem more like one group or three separate individuals?") assessed feelings related to group development, whereas others (e.g., "How much do you think the other members of this work group know about your skills at assembling the radio?") assessed thoughts related to transactive memory. Scores within each set of questions were highly correlated, so they were averaged together, first within and then across subjects, to create two composite indexes for each group. Scores on the group development index were significantly higher in the group training and team-building conditions than in the individual training or reassignment conditions. And scores on the transactive memory index were significantly higher in the group training condition than in the individual training, team-building, or reassignment conditions. The two new conditions thus seemed to affect groups as we hoped they would: The team-building condition encouraged group development, without producing transactive memory systems, while the reassignment condition disabled such systems by making them irrelevant.

Our next set of analyses focused on group performance, measured in the same three ways described earlier. As in the first experiment, we found that groups trained in different ways assembled their radios at about the same speed. There were, however, significant differences among conditions in both procedural recall and assembly errors. These differences are shown in Figures 3.4 and 3.5. As we expected, group training led to better performance outcomes than did any of the other training methods, which did not differ from one another.

Videotapes of the groups were again viewed by two judges, one of whom was blind to the research hypotheses and to each group's condition. As before, these judges were given a list of specific behaviors exemplifying various factors. Three of those factors (memory differentiation, task coordination, and task credibility) reflected the operation of transactive memory systems, whereas the others (task motivation, group cohesion, and social identity) did not. The judges watched each

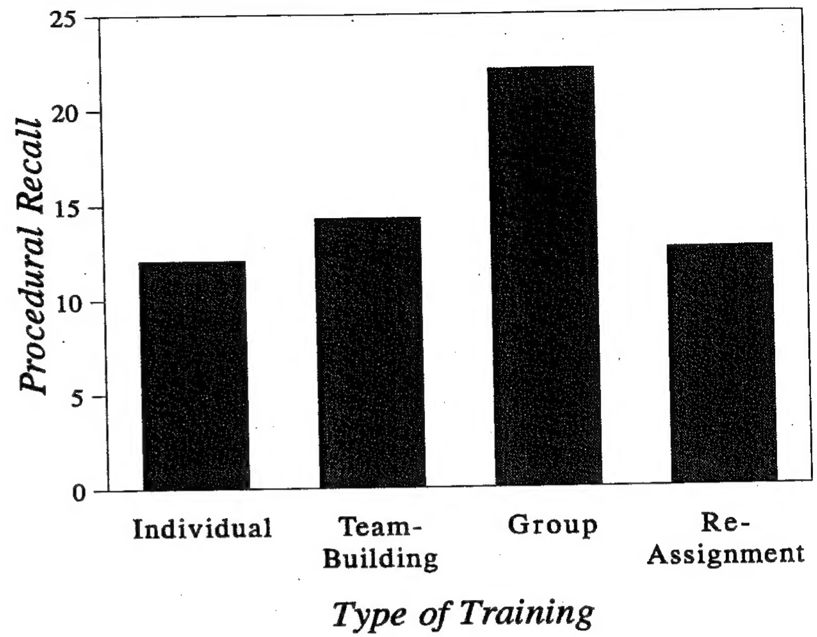

Figure 3.4. Effects of Various Training Methods on Procedural Recall in the Second Experiment

videotape and made an overall rating of the group on each factor. Once again, intraclass correlations indicated that these ratings were made reliably. Scores on the transactive memory factors were highly correlated, so we again created a composite index by averaging each group's scores on those factors. As we expected, scores on that transactive memory index were significantly higher in the group training condition than they were in the other three conditions, which did not differ from one another. Among the remaining factors, only social identity differed significantly across conditions. Social identity scores were higher in the group training and team-building conditions than in the individual training or reassignment conditions.

Were the effects of training methods on group performance mediated by transactive memory systems? Once again, we explored that issue using multiple regression analyses, following the logic of

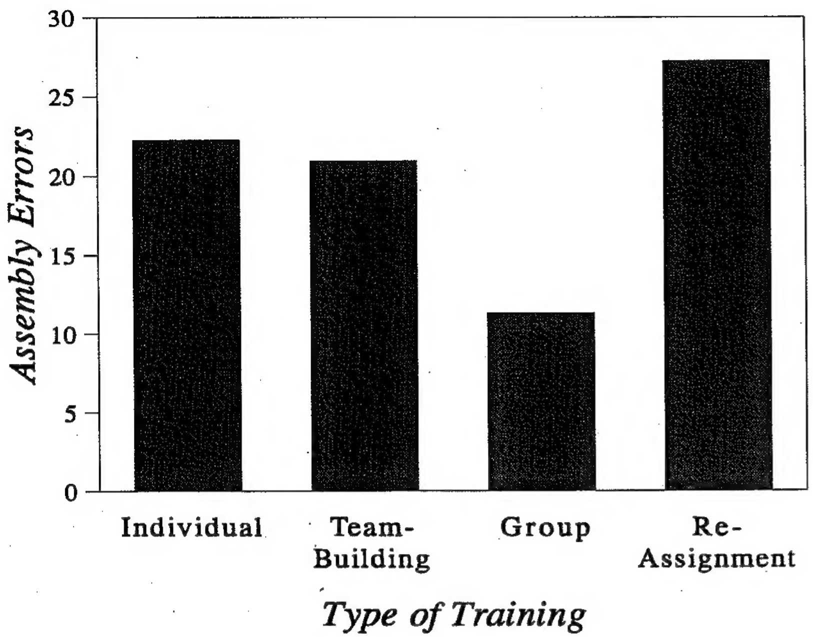

Figure 3.5. Effects of Various Training Methods on Assembly Errors in the Second Experiment

Figure 3.3. Assembly errors served again as our measure of group performance for these analyses, but the four training conditions were now coded as three dummy variables, using a scheme (Cohen & Cohen, 1983) that contrasted the group training condition with the other three conditions. We began, as before, by regressing assembly errors on training methods. The effects of training methods were significant and accounted for 27% of the variance. We then regressed transactive memory on training methods. The effects of training methods were also significant and accounted for 37% of the variance. Finally, we regressed assembly errors on training methods and transactive memory simultaneously. The effects of transactive memory on assembly errors were significant, but the effects of training methods were not. Together, training methods and transactive memory accounted for 47% of the variance in assembly errors.

Social identity could have mediated some of the effects of training methods on group performance, but an analogous set of regression analyses involving that factor revealed no evidence of such mediation. When differences among groups in social identity were taken into account, group training still produced significantly better performance outcomes than any other form of training.

As we noted earlier, this experiment was designed with several goals in mind. One goal, to replicate the results from our first experiment, was clearly achieved. Once again, group training produced better performance than did individual training, and this advantage was due to the development of transactive memory systems. We also wanted to evaluate the contributions of group development and strategic learning to the performance advantages of group training. Those contributions appeared to be minimal, although further research is needed before any final conclusions are drawn. During the training sessions, groups in one condition participated in a team-building exercise that encouraged their development, whereas groups in another condition learned generic strategies for coping with coordination problems. However, few of the groups in either condition had strong transactive memory systems. Groups in the team-building condition lacked the information required for such systems, and groups in the reassignment condition had their systems disabled through membership changes. Without stronger transactive memory systems, neither group development nor strategic learning was sufficient to improve group performance. Groups in both the new conditions performed worse than those in the group training condition and no better than those in the individual training condition. Their poor performance can be attributed to weak transactive memory systems, because training methods no longer affected group performance when differences in the strength of those systems were taken into account. Finally, this experiment also provided a glimpse of how turnover can influence the operation of transactive memory systems. Groups in the reassignment condition experienced sudden and dramatic turnover, which probably created confusion and dismay among their members and thereby harmed their performance. Yet the damage was less serious than it might have been. Although these

groups performed poorly, their performance was no worse than that of groups in the individual training or team-building conditions. Perhaps if work groups were forewarned about turnover (so they could prepare for it), or if fewer group members were lost all at once, then serious disruptions in the operation of transactive memory systems could be avoided. Further research, again, is needed on this issue.

WHAT NOW?

We are currently working on a third experiment, whose primary goal is to prove that group training really does help workers to learn who knows what. In the experiments just described, such knowledge was never measured directly. Videotape measures of several behaviors that reflect the operation of transactive memory systems were examined, but the systems themselves were not assessed. Our third experiment, like the others, involves small groups of college students who are first trained to assemble radios and then tested to see how well they can perform that task. The effects of group and individual training are compared, much as they were in the earlier experiments. However, there is one critical difference: The subjects expect to work in groups during their testing sessions, and groups are in fact scheduled for those sessions, but all of the testing is actually done at the individual rather than the group level. When they arrive for their testing session, the members of each group are first asked to complete questionnaires that measure (in various ways) their beliefs about each person's abilities, skills, and knowledge. Then, working alone, each subject is asked to recall the procedure for assembling a radio and to actually assemble a radio as quickly and accurately as possible. This procedure allows us to assess the actual abilities, skills, and knowledge of every group member, which can then be compared with the subjects' beliefs about such matters, as revealed by the questionnaires. Our general prediction is that more accurate beliefs will be found in groups whose members are trained together rather than apart. The experiment will also allow us to explore other issues, such as whether group training encourages social loafing or free riding (Karau & Williams, 1993; Shepperd, 1993). Group training makes it unnecessary for everyone to learn everything about a task; people can specialize in ways that reflect their talents. But group training also allows some people to learn nothing about a task, if they believe that the group can succeed without them. Although group performance is better when workers are trained together, individual performance may thus be worse.

There are, naturally, many other issues that we plan to explore as our program of research progresses. One set of issues involves the kinds of work experiences that help transactive memory systems to develop. Task experiences are presumably more valuable than social experiences, but even the latter may be helpful if accurate information about relative competence (across workers and tasks) is conveyed. Task experiences may not be necessary at all, if (for example) the manager of a newly formed work group provides its members with information about their abilities, skills, and knowledge, or the workers are led to provide such information themselves. And what happens to work groups whose tasks change—to what extent are transactive memory systems generalizable? Another set of issues involves factors that might moderate the effects of group training on work group performance. Such factors include characteristics of the group, its task, and the work setting (see Liang et al., 1995). For example, training the members of a work group together rather than apart may be especially helpful when (a) the group's composition is heterogeneous, (b) its task requires coordinated effort, or (c) the work setting is stressful. In all of these situations, a transactive memory system is likely to improve group performance, which otherwise might be poor. A final set of issues involves the potential risks of group training. These risks, some of which we have already noted, include social loafing/free riding, the development of group norms that inhibit productivity, and susceptibility to the harmful effects of turnover. Whether the risks of group training ever outweigh its benefits, and how those risks can be minimized or avoided altogether, are questions that we hope to answer.

Our research program focuses on who knows what within work groups, but other cultural features of such groups seem interesting

and important as well. Workers seek collective answers to many questions about themselves, their groups, and the work that they perform (see Figure 3.1). As time passes and shared experiences accumulate, these questions are gradually answered. The resulting knowledge is incorporated into the group's culture and expressed through a variety of customs. Social scientists from several disciplines are now exploring work group cultures, often with intriguing results. Rentsch (1990), for example, has found that different work groups within an organization may interpret the same events in quite different ways. As a result, the impact of a particular event, such as reengineering the organization, can vary considerably from one work group to another. And the role that shared interpretations of events may play in coordinating activities within work groups has been analyzed by Weick and Roberts (1993), who argued that dangerous mistakes can be made when workers ignore or disagree about their task interdependence. Clearly, there is a growing interest in studying work group cultures, not only because of the potential practical value of such research, but also because of the theoretical insights that it may provide.

One of those insights is that cognition can indeed be a group, as well as an individual activity. People often collaborate in the processing of information about matters that concern them. What's surprising is not that such collaboration occurs, but that so few social psychologists (who claim social influence as their domain) acknowledge or investigate it. Social psychology is, ironically, a science that emphasizes individual behavior. Several observers (e.g., Forgas, 1983b; Pepitone, 1981; Steiner, 1986) have commented on this troubling issue, and their comments, which are both thoughtful and emotional, include a variety of predictions about the future. No one can be sure what the future holds, but as we noted earlier, there does seem to be a resurgence of interest among social psychologists in studying small groups, and some of that research involves socially shared cognition. Our own research, along with much of the other research described in this book, shows that group members can process information together, that some of their activities in this regard can be studied scientifically, and that evidence about such activities can clarify important aspects of group life.